

The Asymptotics of Quantum Max-Flow Min-Cut

Matthew B. Hastings

¹*Station Q, Microsoft Research, Santa Barbara, CA 93106-6105, USA*

²*Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA 98052, USA*

The quantum max-flow min-cut conjecture relates the rank of a tensor network to the minimum cut in the case that all tensors in the network are identical¹. This conjecture was shown to be false in Ref. 2 by an explicit counter-example. Here, we show that the conjecture is almost true, in that the ratio of the quantum max-flow to the quantum min-cut converges to 1 as the dimension N of the degrees of freedom on the edges of the network tends to infinity. The proof is based on estimating moments of the singular values of the network. We introduce a generalization of “rainbow diagrams”⁴ to tensor networks to estimate the dominant diagrams. A direct comparison of second and fourth moments lower bounds the ratio of the quantum max-flow to the quantum min-cut by a constant. To show the tighter bound that the ratio tends to 1, we consider higher moments. In addition, we show that the limiting moments as $N \rightarrow \infty$ agree with that in a different ensemble where tensors in the network are chosen independently; this is used to show that the distributions of singular values in the two different ensembles weakly converge to the same limiting distribution. We present also a numerical study of one particular tensor network, which shows a surprising dependence of the rank deficit on $N \bmod 4$ and suggests further conjecture on the limiting behavior of the rank.

PACS numbers:

I. INTRODUCTION

The quantum max-flow min-cut conjecture was introduced in Ref. 1. This conjecture relates the rank of a tensor network for a generic choice of tensor to a maximal classical flow (or minimal cut) on the graph corresponding to the tensor network. In Ref. 2, various forms of the quantum max-flow min-cut conjecture were considered, and the conjectures were in fact shown to be false. Here, we consider a particular version of the conjecture, called version 2 in that paper. Even though the conjecture is not true, we show that the ratio of the actual rank to the rank predicted by the conjecture converges to 1 in the limit of large dimension N of the degrees of freedom on the edges of the network. The proof is statistical in nature, relying on a random choice of tensor in the network in a particular Gaussian ensemble.

We begin by reviewing that particular form of the conjecture (in fact, we consider a special case of the conjecture in that paper, in which all the vertices in the graph have the same degree and all edges have the same capacity; however, our results can be fairly straightforwardly extended to the general case of the conjecture). We slightly modify the notation in that paper.

Consider a tensor network. The tensor network is defined by several pieces of data. First, there is a graph G , with some open edges (i.e., edges that attach to only one vertex inside the graph) and some closed edges (edges which attach to two vertices). Second, for each edge, there is an integer, called the capacity in Ref. 2. Finally, for each vertex, there is a tensor; the number of indices of the tensor is equal to the degree of that vertex and each index of the tensor corresponds to a distinct edge attached to that vertex; each index ranges over a number of possible values equal to the capacity of that edge. The entries of the tensor are complex numbers (one could also consider the case that they are real numbers; this would require some modifications of the techniques here).

By contracting the tensor network, the tensor network assigns a complex number for each choice of indices on the open edges. We partition the open edges into two sets, called the input set and output set. Define D_S to equal the product of the capacities of all input edges and define D_T to equal the product of the capacities of all output edges. This contraction of the tensor network defines a linear operator L from a complex vector space of dimension D_S to one of dimension D_T .

For a given tensor network, the rank of this linear operator is termed the quantum flow of the network. We define two sets, S and T ; these sets are the open ends of the input and output open edges, respectively. We let V be the set of vertices in the graph, not including S and T . We let $\bar{V} = S \cup T \cup V$; below when we refer to vertices, we only mean vertices in V . A cut is a partition of \bar{V} into $\bar{S} \cup \bar{T}$, where $S \subset \bar{S}$ and $T \subset \bar{T}$. The cut set of the cut is the set of edges (v, w) with $v \in \bar{S}$ and $w \in \bar{T}$. For a given graph, and a given cut set separating S from T , define D_C to equal the product of the capacities of the edges in the cut set. Then, define the quantum min-cut of the network to be the minimum of D_C over all cuts.

For a given graph and given capacities, for any choice of tensors, the quantum flow is bounded by the quantum min-cut. We briefly sketch the proof². Given a cut set C , cutting the edges in the cut set separates the graph into two graphs, and similarly the tensor network can be split into two tensor networks, one with input S and output C

and one with input C and output T . Letting L_1 and L_2 be the linear operators defined by these networks, the linear operator L can be written as a product $L = L_2 L_1$. Clearly, since L_1 is a map to a space of dimension D_C , it has rank at most D_C and hence so does L .

This leaves the question: under what circumstances will the quantum flow equal the quantum min-cut? Suppose that all edges have the same capacity. We denote this capacity by N (c was used in Ref. 2; here is one place where we change notation to be more suggestive of “large N ” limits in physics). Then, suppose that all vertices have the same degree d . Choose a single tensor \mathcal{T} which has d indices, each ranging from 1 to N . We then use the *same* tensor \mathcal{T} on all vertices of the graph. For each vertex, choose some ordering assigning indices of the tensor to edges attached to that vertex. Then, for a given d , N , graph, and choice of ordering, define the “quantum max-flow” to be the maximum of the rank over all tensors \mathcal{T} .

Then, the second version of the quantum max-flow min-cut conjecture is that the quantum max-flow is equal to the quantum min-cut. This conjecture was shown to be false.

In fact, the conjecture considered in Ref. 2 is slightly more general, as edges are allowed to have different capacities. Then any two vertices which have the same degree and which have the same sequence of capacities of edges attached to the vertex (with the sequence of edges ordered by the assignment of edges to indices of the tensor) are said to have the same “valence type” and any two vertices with the same valence type have the same tensor. Our results can be extended to this case.

Our main result is that the conjecture is “asymptotically true”, in that as N becomes large, the ratio of the quantum max-flow to the quantum min-cut converges to 1. Write $QMC(G, N)$ to denote the quantum min-cut for a given graph G and given capacity N . Write $QMF(G, N, O)$ to denote the quantum max-flow for a given graph G and ordering O (the symbol L was used for the ordering in Ref. 2 but we use O instead to avoid confusion with L for the linear operator). For a graph G , let $MC(G)$ denote the minimum cut of G , where each edge is assigned capacity 1 to determine the min cut. Then,

$$QMC(G, N) = N^{MC(G)}. \quad (I.1)$$

We show that

Theorem 1. *For all G, O ,*

$$QMF(G, N, O) = QMC(G, N) \cdot (1 - o(1)). \quad (I.2)$$

Here, we use a big-O notation where we consider asymptotic behavior as a function of N . The constant factors hidden by the big-O notation may depend on G, O . Of course, the fact that the flow is bounded by the quantum min-cut implies that $QMF(G, N, O) \leq QMC(G, N)$. The new result will be to lower bound $QMF(G, N, O)$.

Before proving theorem 1, we will prove a weaker theorem:

Theorem 2. *For all G, O ,*

$$QMF(G, N, O) = \Theta(QMC(G, N)). \quad (I.3)$$

The proof of this will rely on estimating the expectation value of $\text{tr}(L^\dagger L)$ and $\text{tr}((L^\dagger L)^2)$, and using a relation between moments of an operator and its rank. Most of the work of the paper will be developing tools to estimate moments.

Before stating the next theorem about moments, we make some definitions:

Definition 1. *Given two tensor networks, $\mathcal{N}_1, \mathcal{N}_2$ with corresponding graphs G_1, G_2 , we define their product to be a tensor network with graph G which is the union of G_1, G_2 . We write the product as a network $\mathcal{N}_1 \cdot \mathcal{N}_2$. The capacity of an edge in G is given by the capacity of the corresponding edge in G_1 or G_2 and the ordering of indices of a vertex in G is given by the ordering of the indices of the corresponding vertex in G_1 or G_2 . If $\mathcal{N}_1, \mathcal{N}_2$ correspond to linear operators L_1, L_2 respectively, then their product corresponds to linear operator $L_1 \otimes L_2$; if $\mathcal{N}_1, \mathcal{N}_2$ both have no open edges, so that the corresponding linear operators are scalars, then the linear operator corresponding to the product is simply the product of these scalars.*

Definition 2. *We say that a tensor network \mathcal{N} with open edges is “connected” if every vertex is connected to some open edge (either input or output) by a path in the graph corresponding to that network.*

Note that a tensor network may be connected and yet the graph corresponding to that network may be disconnected. See Fig. I.1.

Suppose a tensor network \mathcal{N} is not connected, with \mathcal{N} corresponding to graph G and linear operator L . Then, let W be the set of vertices which are not connected to an input or output edge. Then, we can write \mathcal{N} as the product




FIG. I.1: Example of a tensor network that is connected in the sense of definition 2. This network has $|V| = 2$ vertices, labelled by solid circles, and $|E| = 6$ edges, all of which are open. There are $|S| = 3$ input edges, shown at the left of the figure and $|T| = 3$ output edges, shown at the right of the figure.

of two networks, $\mathcal{N}_1, \mathcal{N}_2$, with linear operators L_1, L_2 and graphs G_1, G_2 , where the vertices in \mathcal{N}_2 correspond to the vertices in W and where \mathcal{N}_1 is connected. Then, L_2 is a scalar which is nonzero for generic choice of \mathcal{T} , so that $\text{rank}(L) = \text{rank}(L_1)$. Further, $MC(G) = MC(G_1)$. So, to prove results about the rank of L in terms of $MC(G)$, we can assume, without loss of generality, that \mathcal{N} is connected. So, unless explicitly stated otherwise, all tensor networks with open edges that we consider will be connected and all linear operators L will correspond to connected tensor networks.

In general, we use $|E|$ to denote the number of edges of a tensor network (including open edges) and use $|V|$ to denote the number of vertices. For notational simplicity, throughout the paper we label closed edges by the pair of vertices (v, w) to which they are attached; the results also apply to the case in which there are multiple edges connecting a pair of vertices in which case an additional label must be given to the edge to indicate for which edge it is; for notational simplicity, we do not write this extra label. We will later define a Gaussian ensemble to choose the tensors. For most of the paper, we consider the case that all the tensors in the network are identical, given by some fixed tensor \mathcal{T} drawn from this ensemble. However, we sometimes also consider the case that the tensors in the network are not identical, and instead are chosen independently from the Gaussian ensemble. If we need to distinguish these two cases, we refer to them as the *identical ensemble* and *independent ensemble*. If not otherwise stated, we are referring to the identical ensemble. We use $E[\dots]$ to denote an expectation value in the identical ensemble and $E_{\text{ind}}[\dots]$ to denote an expectation value in the independent ensemble.

We can now state the following theorem proven later.

Theorem 3. Consider a tensor network \mathcal{N} (assumed connected as explained above) with corresponding graph G . Then for the linear operator L corresponding to the tensor network, for any $k \geq 1$, we have

$$E[\text{tr}((L^\dagger L)^k)] = c(G, k) \cdot N^{k|E| - (k-1)MC(G)} + \mathcal{O}(N^{k|E| - (k-1)MC(G) - 1}), \quad (\text{I.4})$$

where $c(G, k)$ denotes a positive integer that depends upon the graph G and on k (the constant $c(G, k)$ does not depend on the ordering O). Further,

$$E_{\text{ind}}[\text{tr}((L^\dagger L)^k)] = c(G, k) \cdot N^{k|E| - (k-1)MC(G)} + \mathcal{O}(N^{k|E| - (k-1)MC(G) - 1}), \quad (\text{I.5})$$

with the same constant $c(G, k)$ in Eq. (I.4) as in Eq. (I.5).

Further, for any graph G , the constant $c(G, k)$ is bounded by an exponential in k , i.e., $C(G, k) \leq c_1 \cdot c_2^k$, where the constants c_1, c_2 depend upon G .

This theorem requires that the network be connected. As a trivial example to show how this is necessary, consider a network which computes a scalar. Suppose that the tensors in the network all are also scalars; i.e., the vertices have degree 0. In this trivial example, suppose that $|V| = 2$. This network is *not* connected. If the two “tensors” at the two different vertices are the scalars x, y , then the operator L is equal to the scalar xy . If one picks x, y independent complex Gaussians with probability distribution function $\frac{1}{\pi} \exp(-|z|^2)$, then one can readily compute $E_{\text{ind}}[\text{tr}(L^\dagger L)] = E_{\text{ind}}[|x|^2 |y|^2] = 1$. However, if we choose x Gaussian and set $y = x$, then $E[|x|^2 |y|^2] = E[|x|^4] = 2$ and so the expectation values would be different in the independent and identical ensembles. This example might seem strange (having degree 0), so the reader can also consider, for example, a tensor network with 2 vertices and d edges, with no open edges so that all edges connect one vertex to the other; using the techniques later to evaluate expectation values, the reader can check that the expectation values will differ.

Thus, if we define

$$K = \left(N^{|E| - MC(G)} \right)^{-1/2} L$$

and define

$$\text{av}(\dots) = \frac{1}{N^{MC(G)}} \text{tr}(\dots),$$

we have

$$E[\text{av}((K^\dagger K)^k)] = c(G, k) + \mathcal{O}(1/N), \quad (\text{I.6})$$

and

$$E_{\text{ind}}[\text{av}((K^\dagger K)^k)] = c(G, k) + \mathcal{O}(1/N). \quad (\text{I.7})$$

The notation $\text{av}(\dots)$ is intended to be suggestive as follows: we know² that if the tensors are chosen independently for generic choice of tensors then L (and hence K) has rank $N^{MC(G)}$. Hence, for given K , the expectation value, of the $2k$ -th moment of a randomly chosen non-zero singular value of a random tensor in the independent ensemble is $E_{\text{ind}}[\text{av}((K^\dagger K)^k)]$. Let μ_N^{ind} be the distribution function of a randomly chosen singular value for a randomly chosen tensor. By the fact that $c(G, k)$ are bounded by an exponential in k , the distributions μ_N^{ind} converge weakly to a limit³. Suppose instead we choose the tensors all identically from the Gaussian ensemble and then randomly choose one of the largest $N^{MC(G)}$ singular values (i.e., if there are $\text{rank}(L) \leq N^{MC(G)}$ non-zero singular values, then with probability $\text{rank}(L)/N^{MC(G)}$ we choose one of the non-zero singular values and otherwise the result we choose a zero singular value). Let this distribution be μ_N . Then, the moments of μ_N are the same as the moments of μ_N^{ind} up to $\mathcal{O}(1/N)$ corrections and so we have the further corollary that

Corollary 1. μ_N converges weakly to the same limit as μ_N^{ind} . The limiting distribution has compact support

Also,

Corollary 2. For all $\epsilon > 0$, for all $n > 0$, there is a constant c such that for all sufficiently large N the probability that the largest singular value of L is greater than or equal to $x\sqrt{N^{|E|-MC(G)+\epsilon}}$ is bounded by cx^{-n} .

Proof. Let λ be the largest singular value of L . $E[\lambda^{2k}] \leq c(G, k)(1 + \mathcal{O}(1/N))N^{k|E|-(k-1)MC(G)}$. So, the probability that $\lambda \geq x\sqrt{N^{|E|-MC(G)+\epsilon}}$ is bounded by

$$\frac{E[\lambda^{2k}]}{x^{2k}N^{k|E|-kMC(G)+k\epsilon}} = c(G, k)x^{-2k}(1 + \mathcal{O}(1/N))N^{MC(G)-k\epsilon}, \quad (\text{I.8})$$

and choosing $k > MC(G)/\epsilon$ and $k \geq n/2$, this is bounded by a constant times x^{-n} for all sufficiently large N . \square

Since if the tensors are chosen independently, one has $QMC(G, N) = QMF(G, N)$ generically, one might naively guess that corollary 1 implies theorem 1 as follows: for independent choice of tensors, the linear operator L will generically have $QMC(G, N)$ nonzero singular values and so one might expect that when the tensors are chosen identically the linear operator L will have nearly $QMC(G, N)$ singular values. The trouble with this naive argument is that it is conceivable that in the independent ensemble the linear operator K will have $QMC(G, N)$ nonzero singular values but that with high probability a constant fraction of them will have magnitude which is $o(1)$ so that μ_N^{ind} will converge to, for example, a sum of a smooth function plus a δ -function at the origin. So, to prove theorem 1 we instead give a more detailed analysis of higher moments.

We begin in section II by lower bounding the rank in terms of traces of moments of the linear operator L . We also define the appropriate Gaussian ensemble for \mathcal{T} in this section, and give a combinatorial method for computing these traces for this ensemble. Then, in section III, we use these methods to estimate the expectation value of $\text{tr}(L^\dagger L)$. In sections IV, V, we show how to estimate expectation values of traces of higher moments of $L^\dagger L$, as well as expectation values of products of such traces; this is done by combining a lower bound in section IV for these expectation values with an upper bound in section V. The techniques for computing the expectation value show that the results are indeed the same, up to $\mathcal{O}(1/N)$, for the two ensembles. In section VI we complete the proof of theorem 1. In section VII we collect some results on variance that are not needed elsewhere but may be of independent interest. In section VIII we present some numerical simulations.

II. MOMENTS BOUND AND DEFINITION OF ENSEMBLE

Let $\text{rank}(L)$ denote the rank of a linear operator L . We have the following bound:

Lemma 1. *For any linear operator L , and any integer $k > 1$,*

$$\text{rank}(L)^{k-1} \geq \frac{\text{tr}(L^\dagger L)^k}{\text{tr}\left((L^\dagger L)^k\right)}, \quad (\text{II.1})$$

assuming that the denominator of the right-hand side is nonzero. In the special case $k = 2$ used later we have

$$\text{rank}(L) \geq \frac{\text{tr}(L^\dagger L)^2}{\text{tr}\left((L^\dagger L)^2\right)}. \quad (\text{II.2})$$

Proof. Let the non-zero singular values of L be $\lambda_1, \dots, \lambda_{\text{rank}(L)}$. Let the vector v be $(\lambda_1^2, \dots, \lambda_{\text{rank}(L)}^2)$. Let the vector w be $(1, 1, \dots, 1)$. Then, by Hölder's inequality applied to vectors v, w ,

$$\sum_{i=1}^{\text{rank}(L)} \lambda_i^2 \leq \left(\sum_{i=1}^{\text{rank}(L)} \lambda_i^{2p} \right)^{1/p} \text{rank}(L)^{1/q} \quad (\text{II.3})$$

for any p, q with $1/p + 1/q = 1$. Choosing $p = k$, $q = k/(k-1)$, and raising the above equation to the k -th power, we get

$$\text{tr}(L^\dagger L)^k \leq \text{tr}\left((L^\dagger L)^k\right) \text{rank}(L)^{k-1}, \quad (\text{II.4})$$

as claimed (in the special case of $k = 2$, we can use Cauchy-Schwarz instead of Hölder). \square

We will prove Theorem 2 by estimating the expected value of the numerator and denominator of Eq. (II.1) for a particular random ensemble of tensors. The ensemble that we choose is that the entries of the tensor \mathcal{T} will be chosen independently and identically distributed, using a Gaussian distribution with probability density

$$\frac{1}{\pi} \exp(-|z|^2)$$

so that $E[|z|^2] = 1$.

We estimate the expectation value of the numerator and denominator of Eq. (II.1) independently, rather than estimating the expectation value of the ratio, and use the following lemma:

Lemma 2. *Let $E[\text{tr}(L^\dagger L)^k]$ and $E[\text{tr}\left((L^\dagger L)^k\right)]$ be given and nonzero for $k > 1$. Then there must exist some tensor \mathcal{T}_0 for which the corresponding linear operator L_0 obeys*

$$\text{rank}(L_0)^{k-1} \geq \frac{E[\text{tr}(L^\dagger L)^k]}{E[\text{tr}\left((L^\dagger L)^k\right)]}. \quad (\text{II.5})$$

Further, since $E[\text{tr}(L^\dagger L)^k] \geq E[\text{tr}\left((L^\dagger L)^k\right)]$,

$$\text{rank}(L_0)^{k-1} \geq \frac{E[\text{tr}\left((L^\dagger L)^k\right)]}{E[\text{tr}\left((L^\dagger L)^k\right)]}. \quad (\text{II.6})$$

Proof. Given $E[\text{tr}(L^\dagger L)^k]$ and $E[\text{tr}\left((L^\dagger L)^k\right)]$, there must exist some \mathcal{T}_0 for which

$$\frac{\text{tr}(L_0^\dagger L_0)^k}{\text{tr}\left((L_0^\dagger L_0)^k\right)} \geq \frac{E[\text{tr}(L^\dagger L)^k]}{E[\text{tr}\left((L^\dagger L)^k\right)]}. \quad (\text{II.7})$$

The result then follows from Eq. (II.1). \square




FIG. II.1: Closed tensor network for computing $\text{tr}(L^\dagger L)$ for the linear operator corresponding to the tensor network shown in Fig. I.1. Closed circles represent tensors \mathcal{T} , while open circles represent tensors $\overline{\mathcal{T}}$. The dashed lines on the 4 open edges are used to indicate that the edges on the right-hand side of the figure should be joined to the edges on the left-hand side of the figure so that the tensor network is closed, computing the trace. Without these dashed lines, the tensor network would have 4 output edges and would compute the linear operator $L^\dagger L$ with the four output edges on the right and the four input edges on the left (while a mathematical expression such as $L^\dagger L$ has its “input on the right” and its “output on the left”, tensor networks are conventionally read from left to right instead).




FIG. II.2: Closed tensor network for computing $\text{tr}((L^\dagger L)^2)$ for the linear operator corresponding to the tensor network shown in Fig. I.1.

For a given tensor \mathcal{T} , traces such as $\text{tr}(L^\dagger L)$ and $\text{tr}(L^\dagger L L^\dagger L)$ are also given by tensor networks; they are tensor networks with no open edges so that their contraction yields a scalar, and this scalar is equal to the desired trace.

Definition 3. We refer to such networks with no open edges as closed tensor networks.

See Figs. II.1, II.2. The notation in Fig. II.1 with closed and open circles to denote \mathcal{T} and $\overline{\mathcal{T}}$ and dashed lines to denote edges that should be joined to compute a trace will be used in figures from here on. We use \mathcal{N} to denote the network with open edges that is used to define L and we use \mathcal{N}_c to denote various different tensor networks with no open edges; the networks \mathcal{N}_c that we consider will correspond to traces such as $\text{tr}(L^\dagger L)$, $\text{tr}(L^\dagger L L^\dagger L)$, and so on; so, \mathcal{N}_c is derived from \mathcal{N} and from the choice of the particular trace.

Definition 4. We introduce notation: $\mathcal{N}_c(\text{tr}(L^\dagger L))$ indicates the closed tensor network corresponding to $\text{tr}(L^\dagger L)$, while $\mathcal{N}_c(\text{tr}(L^\dagger L L^\dagger L))$ indicates the tensor network corresponding to $\text{tr}(L^\dagger L L^\dagger L)$, and so on. Additionally, we consider closed tensor networks which correspond to products of traces, so that $\mathcal{N}_c(\text{tr}(L^\dagger L)\text{tr}(L^\dagger L L^\dagger L))$ indicates the tensor network corresponding to $\text{tr}(L^\dagger L)\text{tr}(L^\dagger L L^\dagger L)$. Such a tensor network $\mathcal{N}_c(\text{tr}(L^\dagger L)\text{tr}(L^\dagger L L^\dagger L))$ is the product of the tensor networks $\mathcal{N}_c(\text{tr}(L^\dagger L))$ and $\mathcal{N}_c(\text{tr}(L^\dagger L L^\dagger L))$.

We now explain how to compute the expectation value of the contraction of a closed tensor network \mathcal{N}_c ; for brevity, we will simply refer to this as “the expectation value of the tensor network”, rather than “the expectation value of the contraction of the tensor network”. The tensor network \mathcal{N}_c is a polynomial in the entries of \mathcal{T} and $\overline{\mathcal{T}}$, where the overline denotes the complex conjugate. Hence, we can use Wick’s theorem to compute the expectation value. Suppose that the tensor network has M different tensors \mathcal{T} and M' different tensors $\overline{\mathcal{T}}$. The expectation value is nonzero only if $M = M'$. Wick’s theorem computes the expectation value by summing over all possible pairings of a tensor \mathcal{T} with a tensor $\overline{\mathcal{T}}$, so that every \mathcal{T} is paired with a unique $\overline{\mathcal{T}}$; that is, there are $M!$ different pairings to sum over. Given a tensor \mathcal{T} with d indices, with entries of the tensor written $\mathcal{T}_{i_1, \dots, i_d}$, we have

$$E[\mathcal{T}_{i_1, \dots, i_d} \overline{\mathcal{T}}_{j_1, \dots, j_d}] = \delta_{i_1, j_1} \dots \delta_{i_d, j_d}. \quad (\text{II.8})$$

These δ -functions can be represented graphically as follows. For each pairing, we define a new graph by removing every vertex (leaving all edges with both ends open) and then for each pair of vertices v, w which are paired with each other, for each $a \in \{1, \dots, d\}$, we attach the end of the edge which corresponded to the a -th index of the tensor at vertex v to the end of the edge which corresponded to the a -th index of the tensor at vertex w . Then, the tensor network breaks into a set of disconnected closed loops. See Fig. II.3.

Definition 5. A closed loop created by a pairing consists of a sequence of edges $(v_1, w_1), (v_2, w_2), \dots, (v_l, w_l)$; no repetition of edges is allowed in a closed loop, but vertices may be repeated. The closed loops created by the pairing




FIG. II.3: A possible pairing of the closed tensor network in Fig. II.1. Dotted lines are used to connect edges which end at vertices which are paired; each vertex in a pair has $d = 3$ different edges: we connected the edges which have the same index for the local ordering. This pairing has 6 different closed loops. There is one other possible pairing for this network; that pairing would have only 3 closed loops.

are such that w_i is paired with v_{i+1} (identifying v_{l+1} with v_1) with the local ordering assigning the same index of the tensor to the edge (v_i, w_i) at w_i as it assigns to edge (v_{i+1}, w_{i+1}) at v_{i+1} . The choice of starting edge in the sequence is irrelevant as is the direction in which the edges are traversed; i.e., the sequences $(v_1, w_1), (v_2, w_2), \dots, (v_l, w_l)$ and $(v_2, w_2), \dots, (v_l, w_l), (v_1, w_1)$ and $(w_l, v_l), \dots, (w_2, v_2), (w_1, v_1)$ all denote the same closed loop.

Suppose for a given pairing π that the number of closed loops is equal to $C(\pi)$. Then, the sum over indices on the edges for the given pairing is $N^{C(\pi)}$. Thus, for a tensor network \mathcal{N}_c with $M = M'$, the expectation value is equal to

$$E[\mathcal{N}_c] = \sum_{\pi \text{ pairing } \mathcal{T} \text{ with } \overline{\mathcal{T}}} N^{C(\pi)}. \quad (\text{II.9})$$

Note that every term in Eq. (II.9) is positive. So, once we have found that there exists some pairing π_0 with some given $C(\pi_0)$, we have established that $E[\mathcal{N}_c] \geq N^{C(\pi_0)}$. Further, the pairing or pairings with the largest $C(\pi_0)$ give the dominant contribution in the large N limit. We define

$$C_{\max} = \max_{\pi} C(\pi), \quad (\text{II.10})$$

and define n_{\max} to be the number of distinct pairings π with $C(\pi) = C_{\max}$; then

$$E[\mathcal{N}_c] = \Theta(N^{C_{\max}}). \quad (\text{II.11})$$

and

$$E[\mathcal{N}_c] = \left(1 - \mathcal{O}(1/N)\right) \cdot n_{\max} N^{C_{\max}}. \quad (\text{II.12})$$

III. ESTIMATING FIRST MOMENT

We now estimate $E[\text{tr}(L^\dagger L)]$.

Lemma 3. *Given a tensor network \mathcal{N} obtained from a graph G with $|V|$ vertices and $|E|$ edges, with corresponding linear operator L , for $\mathcal{N}_c(\text{tr}(L^\dagger L))$ we have*

$$C_{\max} = |E|, \quad (\text{III.1})$$

and

$$n_{\max} = 1. \quad (\text{III.2})$$

Proof. First, we explicitly give a pairing π for which $C(\pi) = |E|$. Note that the number of vertices in the graph G' for tensor network $\mathcal{N}_c(\text{tr}(L^\dagger L))$ equals $2|V|$. There are $|V|$ vertices with tensor \mathcal{T} and $|V|$ vertices with tensor $\overline{\mathcal{T}}$. There is an obvious pairing π of the vertices in G' exemplified in Fig II.3 for one particular network. For this pairing π , we have $C(\pi) = |E|$. We can define this pairing formally for arbitrary G as follows. Suppose G has vertex set \mathcal{V} , and edge set \mathcal{E} . Then, G' has vertices labelled by a pair $(v; \sigma)$ where $v \in \mathcal{V}$ and $\sigma \in \{1, 2\}$. If $\sigma = 1$ then the vertex has tensor \mathcal{T} , and if $\sigma = 2$ then the vertex has tensor $\overline{\mathcal{T}}$. If (v, w) is an edge in G , then $((v; \sigma), (w; \sigma))$ is an edge of G' for $\sigma \in \{1, 2\}$. Additionally, for every open edge in G there is an edge in G' ; for each open edge attached to a vertex v , then there is an edge $((v; 1), (v; 2))$ in G' . There are no other edges in G' other than those given by these rules. The ordering of edges attached to vertices in G' is defined in the obvious way: if an edge e in G attached to a vertex v

corresponds to the j -th index of the tensor, then the edge in G' obtained from e in the above rules attached to $(v; \sigma)$ also corresponds to the j -th index of the tensor. This pairing π is then the pairing which pairs each vertex $(v; 1)$ with $(v; 2)$. This pairing gives one closed loop for every edge of G so that $C(\pi) = |E|$.

We next show that there is no pairing π for which $C(\pi) > |E|$. A closed loop corresponds to a sequence of edges in G' which we write as $(v_1, w_1), (v_2, w_2), \dots, (v_l, w_l)$ for some l ; we say that such a loop “has l edges”. The pairing is such that w_i is paired with v_{i+1} for $i < l$ and w_l is paired with v_1 . If $l = 1$, then v_1 has a tensor \mathcal{T} and w_1 has a tensor $\overline{\mathcal{T}}$ and so this edge in G' is obtained from an open edge in G . If $l > 1$, then vertices v_i and w_i for odd i have tensors \mathcal{T} while for even i they have tensors $\overline{\mathcal{T}}$. So, if $l > 1$, then l is even.

Let n_1 be the number of closed loops in the pairing with $l = 1$; n_1 is bounded by the number of open edges in G , which is equal to $|S| + |T|$. The number of closed loops with $l > 1$ is then bounded by $(1/2)(|E'| - n_1)$, where $|E'|$ is the number of edges in G' , since every closed loop with $l > 1$ must have at least two edges. Note that $|E'| = 2|E| - |S| - |T|$.

So, the number of closed loops is bounded by

$$C(\pi) \leq n_1 + (1/2)(|E'| - n_1), \quad (\text{III.3})$$

which is maximal when n_1 is as large as possible, i.e., $n_1 = |S| + |T|$. In this case, the maximum number of closed loops equals $|E|$. So, $C_{\max} = |E|$.

We now show that $n_{\max} = 1$. Consider a pairing π with $C(\pi) = |E|$ so that $n_1 = |S| + |T|$. Thus, for every open edge in G , there must be a closed loop containing just the edge in G' corresponding to that open edge in G . Further, to have $C(\pi) = |E|$, we must have that no loops have more than two edges, as otherwise $C(\pi) < n_1 + (1/2)(|E'| - n_1)$. A loop with two edges corresponds to a sequence $(v_1, w_1), (v_2, w_2)$ with w_1, v_2 paired and v_1, w_2 paired. This then allows us to show that $n_{\max} = 1$ as follows. Since for every open edge in G , there is a closed loop of length $l = 1$ containing just the corresponding edge in G' , any vertex $(v; 1)$ which is attached to an open edge must be paired with $(v; 2)$. That is, for all vertices $v \in G$ attached to an open edge, we pair $(v; 1)$ and $(v; 2)$. Now consider a vertex $w \in G$ which neighbors some vertex $v \in G$ such that we pair $(v; 1)$ and $(v; 2)$. Then, there is some edge $((w; 1), (v; 1))$ and this edge must be in a closed loop with two edges. Since we pair $(v; 1)$ with $(v; 2)$, there must be a loop containing edges $((w; 1), (v; 1))$ and $((v; 2), (w; 2))$. Then, since this loop must have length 2, we must pair $(w; 1)$ with $(w; 2)$. Let P be the set of vertices $v \in G$ such that we pair $(v; 1)$ with $(v; 2)$; we have shown that P contains all vertices attached to an open edge and P contains all vertices connected to a vertex in P by an edge. So, since the network is connected, P must contain all vertices and so there is indeed only one such pairing. \square

IV. LOWER BOUND FOR HIGHER MOMENTS

We now lower bound $E[\text{tr}(L^\dagger L L^\dagger L)]$ and other higher moments. First let G' denote the graph for tensor network $\mathcal{N}_c(\text{tr}(L^\dagger L L^\dagger L))$. Let us formally define G' , in a way similar to that in which a different graph (also called G') was defined in lemma 3. Suppose G has vertex set \mathcal{V} , and edge set \mathcal{E} . Then, G' has vertices labelled by a pair $(v; \sigma)$ where $v \in \mathcal{V}$ and $\sigma \in \{1, 2, 3, 4\}$. If $\sigma \in \{1, 3\}$ then the vertex has tensor \mathcal{T} , and if $\sigma \in \{2, 4\}$ then the vertex has tensor $\overline{\mathcal{T}}$. If (v, w) is an edge in G , then $((v; \sigma), (w; \sigma))$ is an edge of G' for $\sigma \in \{1, 2, 3, 4\}$. Additionally, for every open edge in G there are two edges in G' ; for each open edge attached to a vertex v , if the edge is an input edge then there are edges $((v; 1), (v; 2))$ and $((v; 3), (v; 4))$ in G' while if the edge is an output edge then there are edges $((v; 2), (v; 3))$ and $((v; 4), (v; 1))$. There are no other edges in G' other than those given by these rules. The ordering of edges attached to vertices in G' is defined in the obvious way: if an edge e in G attached to a vertex v corresponds to the j -th index of the tensor, then the edge in G' obtained from e in the above rules attached to $(v; \sigma)$ also corresponds to the j -th index of the tensor. If we instead consider a higher moment $E[\text{tr}((L^\dagger L)^k)]$, we define a graph G' for tensor network $\mathcal{N}_c(\text{tr}((L^\dagger L)^k))$ similarly. Now G' has vertices labelled by a pair $(v; \sigma)$ where $v \in \mathcal{V}$ and $\sigma \in \{1, 2, \dots, 2k\}$. If σ is odd then the vertex has tensor \mathcal{T} , and if σ is even then the vertex has tensor $\overline{\mathcal{T}}$. If (v, w) is an edge in G , then $((v; \sigma), (w; \sigma))$ is an edge of G' for all σ . Additionally, for every open edge in G there are k edges in G' ; for each open edge attached to a vertex v , if the edge is an input edge then there are edges $((v; \sigma), (v; \sigma + 1))$ for σ odd while if the edge is an output edge then there are edges $((v; \sigma), (v; \sigma + 1))$ for σ even. We regard σ as periodic mod $2k$, so that $\sigma = 2k + 1$ is the same as $\sigma = 1$. The ordering of edges attached to vertices in G' is defined in the obvious way: if an edge e in G attached to a vertex v corresponds to the j -th index of the tensor, then the edge in G' obtained from e in the above rules attached to $(v; \sigma)$ also corresponds to the j -th index of the tensor.

Thus, for example, for the closed tensor network in Fig. II.2, the four vertices in the top row correspond to $\sigma = 1, 2, 3, 4$ from *left to right* (recall that the input of the tensor network is on the left), as do the four vertices in the bottom row.

We now show the lower bound




FIG. IV.1: Example of the pairing defined in lemma 4. The network is the same as in Fig. II.2. The thin curving vertical lines represent minimum cuts of the network. There are four such lines, each cutting one of the cases $\sigma = 1, 2, 3, 4$. The lines for odd σ are reflected in the horizontal direction compared to those for even σ . The vertical dotted lines represent reflection planes: reflecting the region between any two neighboring vertical curved lines about one of these vertical dashed lines leaves the region unchanged except for interchanging tensors \mathcal{T} and $\bar{\mathcal{T}}$ (we have drawn one of these lines at the right-hand side of the figure, indicating a reflection relating vertices at the right-hand side to those at the left-hand side). The pairing in lemma 4 pairs vertices related by such a reflection.

Lemma 4. *Let linear operator L correspond to a tensor network \mathcal{N} obtained from a graph G with $|V|$ vertices and $|E|$ edges. Then, for $\mathcal{N}_c(\text{tr}((L^\dagger L)^k))$ for $k \geq 1$ we have*

$$C_{\max} \geq k|E| - (k-1)MC(G). \quad (\text{IV.1})$$

Proof. The case $k = 1$ is already given above. So, assume $k > 1$.

Suppose the lemma does not hold, so that $C_{\max} < k|E| - (k-1)MC(G)$ and hence $C_{\max} \leq k|E| - (k-1)MC(G) - 1$. Then from lemma 3, for sufficiently large N we would have

$$\frac{E[\text{tr}(L^\dagger L)^k]}{E[\text{tr}((L^\dagger L)^k)]} \geq c \cdot N^{(k-1)MC(G)+1}, \quad (\text{IV.2})$$

for some positive constant c (the ratio of expectation values would asymptotically tend to $1/n_{\max}$, where here n_{\max} refers to the number of pairings with maximal $C(\pi)$ for tensor network $\mathcal{N}_c(\text{tr}((L^\dagger L)^k))$). Then, from lemma 2, for some choice of tensor \mathcal{T}_0 the corresponding linear operator L_0 obeys $\text{rank}(L_0) \geq c \cdot N^{MC(G)+1/(k-1)}$, which is asymptotically larger than $QMC(G, N)$, contradicting the fact that $\text{rank}(L) \leq QMC(G, N)$.

In addition to this proof, we give an alternative proof by explicitly giving a pairing π for which $C(\pi) = k|E| - (k-1)MC(G)$. We exemplify this pairing for a particular network in Fig. IV.1. Consider a minimum cut, with corresponding sets \bar{S}, \bar{T} . For $v \in \bar{T}$, we pair $(v; \sigma)$ with $(v; \sigma + 1)$ for odd σ , while for $v \in \bar{S}$ we pair $(v; \sigma)$ with $(v; \sigma + 1)$ for even σ . Then, for every edge (v, w) for vertices $v, w \in \bar{T} \setminus T$ there are k closed loops, corresponding to edges $((v; \sigma), (w; \sigma))$ and $((w; \sigma + 1), (v; \sigma + 1))$ for odd σ ; similarly, for every edge (v, w) for vertices $v, w \in \bar{S} \setminus S$ there are k closed loops, corresponding to edges $((v; \sigma), (w; \sigma))$ and $((w; \sigma + 1), (v; \sigma + 1))$ for even σ . These edges (v, w) for $v, w \in \bar{T} \setminus T$ or $v, w \in \bar{S} \setminus S$ are not in the cut set. Similarly, for every edge in the output set or in the input set (assuming that the edge is not in the cut set) there are k closed loops. However, for each edge (v, w) in the cut set, there is only one closed loop, corresponding to edges $((v; 1), (w; 1)), ((w; 2), (v; 2)), \dots, ((w; 2k), (v; 2k))$.

Thus, the number of closed loops is equal to k times the number of edges not in the cut set, plus the number of edges in the cut set, which equals $k|E|$ minus $k-1$ times the number of edges in the cut set. \square

V. UPPER BOUND FOR HIGHER MOMENTS AND ITS REALIZATION BY “DIRECT PAIRINGS”

The pairing in the lemma 4 has a certain structure, which we now define (that pairing is not the only pairing consistent with this structure).

Definition 6. *Consider an arbitrary closed tensor network, \mathcal{N}_c , with some vertices having tensor \mathcal{T} and some having tensor $\bar{\mathcal{T}}$. Suppose that there are two vertices, v, w with v having tensor \mathcal{T} and w having tensor $\bar{\mathcal{T}}$. Further, suppose that there is at least one edge from v to w which has the property that the local ordering of indices makes that edge*




FIG. V.1: Example of a direct pairing that does not have maximal $C(\pi)$. The notation is the same as in Fig. IV.1, except that the thin curving vertical lines are cuts which are *not* min cuts.

correspond to the same index for tensor \mathcal{T} as it does for tensor $\overline{\mathcal{T}}$. Then, define a new closed tensor network, \mathcal{N}'_c by removing the vertices v, w from \mathcal{N}_c ; every edge from v to w is removed, while for every pair of edges (u, v) and (w, x) for which the local ordering gives the same tensor index at v as at w , we replace that pair with a single edge (u, x) , defining the local ordering in the obvious way, so that the edge (u, x) corresponds to the index of the tensor at u that (u, v) did and corresponds to the index of the tensor at x that (w, x) did. Then, we say that a network \mathcal{N}'_c constructed in this fashion is a “one-step direct subnetwork” of \mathcal{N}_c ; more specifically, we call it a “one-step direct subnetwork made by pairing v, w in \mathcal{N} ” to indicate how it is constructed. Any network \mathcal{N}'_c constructed by zero or more steps of the above procedure is termed a “direct subnetwork” of \mathcal{N}_c .

Given a sequence of direct subnetworks, $\mathcal{N}_c(1)$ pairing $(v(1), w(1))$ in \mathcal{N}_c and $\mathcal{N}_c(2)$ pairing $(v(2), w(2))$ in $\mathcal{N}_c(1)$, and so on, we define a partial pairing of the original network \mathcal{N}_c . This partial pairing pairs $v(1)$ with $w(1)$, pairs $v(2)$ with $w(2)$, and so on. This is a partial pairing as it may pair only a subset of the vertices. Such a partial pairing is termed a direct partial pairing. If all vertices are paired (so that the last direct subnetwork in the sequence has no vertices), then this is termed a direct pairing.

For each direct partial pairing there are several possible sequences of direct subnetworks $\mathcal{N}_c(1), \dots$, but the last network in the sequence is unique, and we say that this is the direct subnetwork determined by that direct partial pairing.

One way to understand direct pairings is to consider the case that the graph has degree $d = 2$ and has one input and one output edge. Then, if the graph has a single vertex, the problem of the singular values of L reduces to a well-studied problem in random matrix theory, studying the singular values of a random square matrix with independent entries. This is the so-called chiral Gaussian Unitary Ensemble^{5,6}. In this case, it is well-known that the dominant diagrams in the large N limit for any moment are the so-called “rainbow diagrams” (these are also called “planar diagrams”)⁴. Even if there is one input and one output edge, but more than one vertex (so that we now study the singular values of a power of a random matrix) the dominant diagrams are still rainbow diagrams. These rainbow diagrams are precisely the direct pairings in this case. If we still stick to the case $d = 2$, with $|S| = |T| = MC(G)$, but allow $|MC(G)| > 1$, dominant diagrams can still be understood as rainbow diagrams: for each of the $MC(G)$ distinct paths from input to output, we draw a rainbow diagram, pairing only vertices which are both in the same path (i.e., given a power $k \geq 1$ so that we have k copies of a given path with tensors \mathcal{T} and k copies with tensors $\overline{\mathcal{T}}$, we pair vertices between copies of that path, but only within a given path, not between paths). Again, these are the direct pairings.

Later, we will use this understanding in terms of rainbow diagrams to better understand the case $d > 2$. Suppose $d > 2$ and we have a direct pairing. We can construct $MC(G)$ edge disjoint paths from input to output. We will show that the direct pairing has the properties that if we consider only the vertices and edges in one of these paths, the result is a rainbow diagram. Since the edge-disjoint paths might share vertices (if $d \geq 4$) this can impose some relationship between the different rainbow diagrams for each paths: i.e., it is not the case that we can choose a distinct rainbow diagram for each path independently. Further, while every pairing that gives rainbow diagrams when restricted to each of these paths will be a direct pairing, not every such direct pairing will have maximal $C(\pi)$; see Fig. V.1.

Definition 7. Consider an arbitrary closed tensor network \mathcal{N}_c , and a one-step direct subnetwork \mathcal{N}'_c made by pairing vertices v, w in \mathcal{N}_c . Let π' be a partial pairing of \mathcal{N}'_c . Then, we say that π' induces a partial pairing of \mathcal{N}_c which pairs v with w and pairs all other vertices as they are paired in π' (every vertex in \mathcal{N}_c other than v, w corresponds to a vertex in \mathcal{N}'_c). Inductively, for any direct partial pairing θ defining a direct subnetwork \mathcal{N}'_c and any pairing π' of \mathcal{N}'_c we define a partial pairing on \mathcal{N}_c induced by π' . We write $\Pi(\pi', \theta)$ to denote this partial pairing.

Definition 8. Given a partial pairing π of a network, we define $C(\pi)$ the number of closed loops created by the partial




FIG. V.2: Example of paths $Q(i)$ as constructed in lemma 8. The two thickened solid lines each represent such a path. For this network, the choice of paths $P(i)$ is non-unique.

pairing in the obvious way: it is the number of distinct closed loops, where as in Definition 5 each closed loop contains edges $(v_1, w_1), \dots, (v_l, w_l)$ such that w_i is paired with v_{i+1} (identifying v_{l+1} with v_1) with the local ordering assigning the same index of the tensor to the edge (v_i, w_i) at w_i . Note that for a partial pairing, some edges might not be in a closed loop.

Lemma 5. Consider direct partial pairing θ determining direct subnetwork \mathcal{N}'_c and let π' be a pairing of \mathcal{N}'_c . Then

$$C(\Pi(\pi', \theta)) = C(\pi') + C(\theta). \quad (\text{V.1})$$

In the special case that \mathcal{N}'_c is a one-step direct subnetwork, $C(\theta) \leq N_E(v, w)$, where $N_E(v, w)$ is the number of edges from v to w , with v, w as in definition 6; this is an equality if the ordering is such that all of these edges correspond to the same index at v as they do at w .

When applying this lemma 5, we will say below that the pairing of vertices v, w “creates $N_E(v, w)$ closed loops” or that the pairing θ “creates $C(\theta)$ closed loops”.

Lemma 6. Let π be a pairing of a closed tensor network \mathcal{N}_c . Assume that there is no one-step direct subnetwork \mathcal{N}'_c with pairing π' such that π is induced by π' . Then,

$$C(\pi) \leq N_E(\mathcal{N}_c)/2, \quad (\text{V.2})$$

where $N_E(\mathcal{N}_c)$ is the number of edges in tensor network \mathcal{N}_c .

Proof. Every closed loop for pairing π must be composed of at least two edges (if it is composed of one edge, then pairing those vertices defines a one-step direct subnetwork). \square

Lemma 7. Let $\mathcal{N}_c = \mathcal{N}_c(\text{tr}((L^\dagger L)^k))$. Let G be the graph corresponding to this tensor network. Let \mathcal{N}'_c be a direct subnetwork of \mathcal{N}_c . Let π' be the partial pairing defined by \mathcal{N}'_c (if \mathcal{N}'_c consists only of closed loops, then π' is a pairing). Then, labelling the vertices of \mathcal{N}_c by pairs $(v; \sigma)$ with $\sigma \in \{1, \dots, 2k\}$ as above, the partial pairing π only pairs $(v; \sigma)$ with $(w; \tau)$ for $v = w$ and σ odd and τ even.

Further, (*) if a pair of vertices $(x; \mu)$ and $(y; \nu)$ in \mathcal{N}'_c are connected by an edge in \mathcal{N}'_c , then either $x = y$ and $\mu \neq \nu \pmod 2$ and the local ordering is such that the edge is assigned the same index of the tensor at $(x; \mu)$ as it is at $(y; \nu)$, or $\mu = \nu \pmod 2$ and there is an edge (x, y) in G and the local orderings agree: if the ordering assigns the j -th index of the tensor to edge (x, y) at x it also assigns the j -th index of the tensor to the edge $((x; \mu), (y; \nu))$ at $(x; \mu)$ and similarly for y and $(y; \nu)$.

Proof. The fact that $\sigma; \tau$ have different parity mod 2 follows because the odd and even vertices correspond to $\mathcal{T}^\dagger, \mathcal{T}$ respectively.

The statement (*) in the last paragraph of the claim of the lemma can be established inductively. The base case is that $\mathcal{N}'_c = \mathcal{N}_c$, where (*) follows trivially. If \mathcal{N}'_c is a direct subnetwork obeying(*) and \mathcal{N}''_c is a one-step direct subnetwork, one can check case-by-case that \mathcal{N}''_c obeys (*).

Once (*) is established, the fact that the pairing only pairs $(v; \sigma)$ with $(w; \tau)$ for $v = w$ and σ odd and τ even follows inductively since one only pairs vertices connected by an edge. \square

Lemma 8. Let $\mathcal{N}_c = \mathcal{N}_c(\text{tr}((L^\dagger L)^k))$. Then, for any direct pairing π ,

$$C(\pi) \leq k|E| - (k-1)MC(G). \quad (\text{V.3})$$

For any pairing π that is not a direct pairing,

$$C(\pi) \leq k|E| - (k-1)MC(G) - 1. \quad (\text{V.4})$$

Finally, for any graph G , n_{\max} is bounded by an exponentially growing function of k .

Proof. First consider the case that π is a direct pairing. Consider a maximal flow on the graph G , giving each edge of the graph capacity 1. By the max-flow/min-cut theorem, the flow is equal to the min cut^{8,9}. Further, the flow on each edge in a maximal flow can be chosen to be an integer, and hence equals 0 or 1 on every edge. Consider the set of edges on which the flow equals 1. This set defines $MC(G)$ edge-disjoint paths from input to output. Call these paths $P(1), P(2), \dots, P(MC(G))$.

Let $v_{i,1}, v_{i,2}, \dots, v_{i,l(i)}$ denote the sequence of vertices in $P(i)$, as the path is traversed from input to output, where $l(i)$ is the total number of vertices in the path $P(i)$. Each path in the graph defines a closed path $Q(i)$ in the network: the path is traversed backwards for each network corresponding to L^\dagger and forward for each network corresponding to L so that it forms a closed path traversing $kl(i)$ vertices; i.e., the path $Q(i)$ traverses vertices $(v_{i,1}; 1), (v_{i,2}; 1), \dots, (v_{i,l(i)}; 1), (v_{i,l(i)}; 2), \dots, (v_{i,2}; 2), (v_{i,1}; 2), \dots$. See Fig. V.2 for an example. If $P(i)$ has e_i edges in the graph, including 1 input edge and 1 output edge, then the number of edges E_i in $Q(i)$ is equal to $2k(e_i - 1)$.

By lemma 7, for a direct pairing, for every such path, vertex $(v_{i,b}; \sigma)$ is paired with $(v_{i,b}; \sigma')$ for some σ' . So, π pairs vertices in $Q(i)$ with other vertices in $Q(i)$. As noted above, in the case of a graph with degree $d = 2$, a direct pairing gives a rainbow diagram. Here, if we consider the subgraph containing only vertices in $Q(i)$, we have a graph with degree 2 and again pairing π defines a “rainbow diagram”. Hence, the number of closed loops formed by edges on the path is equal to

$$\frac{1}{2}E_i + 1 = ke_i - (k - 1).$$

(One can derive this number of closed loops using the known result for rainbow diagrams, or one can also derive it inductively: given a graph with degree 2 containing 1 closed loop with $e_i > 2$ edges, pairing two vertices v, w gives a one-step direct subnetwork with 2 fewer edges and creates $N_E(v, w) = 1$ edges between them, while for $e_i = 2$, the pairing gives two closed loops.) Hence, the total number of closed loops formed in all such paths in the network is equal to

$$k \sum_{i=1}^{MC(G)} e_i - (k - 1)MC(G).$$

Let Q_E be the set of edges in \mathcal{N}_c which are in some path $Q(i)$ and Q_V be the set of vertices in \mathcal{N}_c which are in some path $Q(i)$. We now count the number of closed loops which do not contain any edges in Q_E ; note that for a direct pairing, every closed loop either contains only edges from Q_E or contains no edge in Q_E . Let P^\perp be the set of edges in G which are not in a path $P(i)$, so $|P^\perp| = |E| = \sum_i e_i$. Note that $S - MC(G)$ of the edges in P^\perp are input edges and $T - MC(G)$ of such edges are output edges so that there are $|P^\perp| - (S - MC(G)) - (T - MC(G))$ edges in G which are not in a path $P(i)$ and which are closed edges. Every edge in P^\perp which is closed corresponds to $2k$ edges in the network, while every edge in P^\perp which is open corresponds to only k edges in the network. The number of closed loops which can be formed by these edges is at most $1/2$ the number of edges which connect \mathcal{T} to \mathcal{T} (i.e., corresponding to the closed edges in G) or \mathcal{T} to \mathcal{T} (i.e., also corresponding to closed edges), plus the number of edges connecting \mathcal{T} to \mathcal{T} (i.e., corresponding to the open edges in G) so that the total number of such closed loops is at most $(1/2)(2k)(|P^\perp| - (S - MC(G)) - (T - MC(G))) + k(S - MC(G)) + (T - MC(G)) = k|P^\perp|$. Hence, the total number of closed loops including edges in Q_E and edges not in Q_E is at most

$$\sum_i e_i + (k - 1)MC(G) + k|P^\perp| = k|E| - (k - 1)MC(G). \quad (V.5)$$

Now, suppose that π is not a direct pairing. We show that $C(\pi) \leq k|E| - (k_i - 1)MC(G) - 1$. Before giving the proof, we give some motivation. The basic idea is that if we consider the edges not in Q_E then no pairing can do better than the direct pairing: the direct pairing has one loop of length 1 for each edge in Q_E corresponding to an output edge in G and has one loop of length 2 for each edge in Q_E not corresponding to an output edge in G , and that is the shortest such a loop can be. On the other hand, when we consider the edges in Q_E , we know that a rainbow diagram gives the optimal pairing for a network with degree $d = 2$ and so any other pairing must be worse. To do this in detail, we will use lemma 6.

The θ be a direct partial pairing of \mathcal{N}_c determining direct subnetwork \mathcal{N}'_c such that π is induced by a pairing π' of \mathcal{N}'_c (i.e., $\pi = \Pi(\pi', \theta)$) and such that there is no one-step direct subnetwork \mathcal{N}''_c of \mathcal{N}'_c with pairing π'' such that π' is induced by π'' . By lemma 6, $C(\pi') \leq N_E(\mathcal{N}'_c)/2$, so by lemma 5, $C(\pi) \leq N_E(\mathcal{N}'_c)/2 + C(\theta)$. We wish to estimate $C(\theta)$ and to estimate $N_E(\mathcal{N}_c) - N_E(\mathcal{N}'_c)$. Recall that when we define a one-step direct subnetwork, the number of edges changes for two reasons: we remove $N_E(v, w)$ edges (and create $N_E(v, w)$ loops) but we also combine other pairs of edges into a single edge.

Let C be the set of these closed loops created by θ . Let C_E be the set of edges in a loop in C . Each closed loop in C contains either only edges in Q_E or contains no edges in Q_E . Consider the loops in C containing edges not in

Q_E . Let R_o be the set of edges in these loops which correspond to output edges in G and let R_i be the set of edges in these loops which do not correspond to output edges in G so that the number of loops in C containing edges not in Q_E is bounded by $|R_o| + (1/2)|R_i|$. f

Consider the loops in C containing edges in Q_E . Each such loop contains edges from at most one path $Q(i)$. In such a path, there are $2k(e_i - 1)$ total edges in \mathcal{N}_c . For each path $Q(i)$ we define a path in \mathcal{N}'_c in the obvious way, so that the path in \mathcal{N}'_c consists of the vertices in \mathcal{N}'_c which are in the path in \mathcal{N}_c . If all edges in $Q(i)$ are in a loop in C then there are $k(e_i - 1) + 1$ closed loops in C containing edges in $Q(i)$. On the other hand, if not all edges in $Q(i)$ are in such a loop, then there are less than $k(e_i - 1) + 1$ closed loops and the number of closed loops is equal to $(1/2)(E_i - E'_i)$, where E_i is the number of edges in $Q(i)$ and E'_i is the number of edges in the corresponding path in \mathcal{N}'_c . So,

$$N_E(\mathcal{N}_c) - N_E(\mathcal{N}'_c) \geq |R_o| + |R_i| + \sum_i (E_i - E'_i). \quad (\text{V.6})$$

Let q be the number of paths $Q(i)$ such that π' pairs all vertices in $Q(i)$. So,

$$\begin{aligned} C(\theta) &= |R_o| + (1/2)|R_i| + (1/2) \sum_i (E_i - E'_i) + q \\ &\leq \frac{N_E(\mathcal{N}_c) - N_E(\mathcal{N}'_c)}{2} + |R_o|/2 + q. \end{aligned} \quad (\text{V.7})$$

So,

$$N_E(\mathcal{N}'_c)/2 + C(\theta) \leq N_E(\mathcal{N}_c)/2 + |R_o|/2 + q. \quad (\text{V.8})$$

The number of open edges in G is $|S| + |T|$ of which $|S| + |T| - 2MC(G)$ edges are not in a path $Q(i)$, so $|R_o| \leq k(|S| + |T| - 2MC(G))$. So,

$$\begin{aligned} N_E(\mathcal{N}'_c)/2 + C(\theta) &\leq N_E(\mathcal{N}_c)/2 + k(|S| + |T|)/2 - kMC(G) + q \\ &= k|E| - kMC(G) + q. \end{aligned} \quad (\text{V.9})$$

So, unless $q = MC(G)$, we have established the desired bound on $C(\pi)$.

So, suppose that $q = MC(G)$ and suppose that $C(\pi) = C_{\max}$. We will show that π is a direct pairing. We will use the assumption that the network is connected. Since $C(\pi) = C_{\max}$, every edge in \mathcal{N}_c which is not in a loop created by θ must either be in a loop of length 1 (if corresponds to an open edge in G) or to a loop of length 2 (otherwise). Consider a vertex $(v; \sigma)$ which is not paired by θ and which is attached to an open edge. Since this edge must be in a closed loop of length 1, $(v; \sigma)$ must be paired with $\sigma \pm 1$ depending on $\sigma \bmod 2$ and on whether it is an input edge or an output edge. Thus, we can pair $(v; \sigma)$ with $(v; \sigma \pm 1)$ in \mathcal{N}'_c giving a further one-step direct subnetwork \mathcal{N}''_c such that a pairing π'' on \mathcal{N}''_c induces π' . So, we can assume that no such vertices exist. Consider instead a vertex $(v; \sigma)$ which is not paired by θ and which neighbors a vertex $(w; \sigma)$ which is paired by θ . The pairing θ is a direct partial pairing so it pairs $(w; \sigma)$ with $(w; \tau)$ for some τ . Then, since the edge $((v; \sigma), (w; \sigma))$ must be in a loop of length 2 in π , the pairing π must pair $(v; \sigma)$ with $(v; \tau)$. So, pairing these vertices would define a one-step direct subnetwork \mathcal{N}''_c . So, there can be no vertices not paired by θ which are attached to output edges or which neighbor a vertex paired by θ ; so, since the network is connected, all vertices are paired by θ and π is a direct pairing.

We now bound n_{\max} to bound the $c(G, k)$. Let π be a direct pairing with $C(\pi) = C_{\max}$. In each path $Q(i)$, the direct pairing must define a rainbow diagram. Further, since every edge not in Q_E is either in a loop of length 1 (if it is an open edge) or a loop of length 2 (otherwise), the pairing of the vertices in Q_V fully determines π . So, we can bound n_{\max} by bounding the number of possible pairings of the vertices in Q_V . For each path $Q(i)$, the pairings of the vertices in that path define some rainbow diagram. If $P(i)$ has $l(i)$ vertices, then $Q(i)$ has $2kl(i)$ vertices. There is some restriction on the pairing of these vertices, as one can only pair vertices corresponding to \mathcal{T} to those corresponding to $\overline{\mathcal{T}}$. Ignoring this restriction to obtain an upper bound, we ask for the number of possible rainbow diagrams pairing $2kl(i)$ vertices (this is a problem that arises in estimating the $2kl(i)$ -th moment in the Gaussian Orthogonal Ensemble where one consider random real symmetric matrices). The number of such rainbow diagrams is at most exponential⁷ in $2kl(i)$ and so the product over paths $Q(i)$ of the number of such rainbow diagrams for each path is at most exponential in $2k \sum_i l(i) \leq 2k|V|$, showing the desired result. The number n_{\max} may be less than the product of these for two reasons: first, not all direct pairings have $C(\pi) = C_{\max}$ and second, if two paths share a vertex then the imposes some relation between the pairings on those two paths. \square

Proof of Theorem 2 By lemmas 4,8, for $\mathcal{N}_c = \mathcal{N}_c(\text{tr}((L^\dagger L)^k))$, we have $C_{\max} = k|E| - (k-1)MC(G)$. This implies Eq. (I.4) in theorem 3 with $c(G, k) = n_{\max}$ (we show below that $c(G, k)$ does not depend on O ; we will not need that

to prove theorem 2). Using Eq. (I.4) for $k = 2$ to estimate $E[\text{tr}((L^\dagger L)^k)]$ and using lemma 3 to estimate $E[\text{tr}(L^\dagger L)]$, we find that

$$\frac{E[\text{tr}(L^\dagger L)]^2}{E[\text{tr}((L^\dagger L)^2)]} \geq \frac{1}{c(G, 2)} N^{MC(G)} - \mathcal{O}(1/N). \quad (\text{V.10})$$

So, by lemma 2, theorem 2 follows.

Proof of Theorem 3 We have shown Eq. (I.4) in theorem 3. To show that Eq. (I.5) holds for the independent ensemble, note that in that ensemble, the only pairings allowed are those in which we pair $(v; \sigma)$ with $(w; \tau)$ for $v = w$. However, by lemma 7 all direct pairings have that property and by lemma 8 the only pairings π with $C(\pi) = C_{\max}$ are direct pairings. To show that $c(G, k)$ indeed does not depend on O , one can either note that the possible direct pairings do not depend on O or one can note that in the ensemble in which tensors are chosen independently, one can freely change the ordering at any vertex without altering the expectation values. The bound on the $c(G, k)$ follows from the bound on n_{\max} in lemma 8.

VI. PROOF OF THEOREM 1

We now prove theorem 1. Let $f(x)$ be some smooth function defined for $0 \leq x < \infty$ with $0 \leq f(x) \leq 1$, $f(x) = 0$ for $x \geq 2$, and $f(x) = 1$ for $0 \leq x \leq 1$ and let $f(x)$ decrease monotonically with increasing x . Let $P_N(\epsilon) = \int f(x/\epsilon) d\mu_N$, with ϵ chosen later. We choose $f(x)$ to be smooth so that as $N \rightarrow \infty$, $P_N(\epsilon)$ converges to some limit $Pr(\epsilon, G) = \int f(x/\epsilon) d\mu$, where we explicitly put the dependence on G in parentheses as we will deal with this probability for several different graphs below. We have

$$\int f(x/\epsilon) d\mu_N \geq 1 - \frac{E[\text{rank}(L)]}{QMC(G, N)}, \quad (\text{VI.1})$$

for all $\epsilon > 0$. So, $Pr(\epsilon, G) \geq \liminf_{N \rightarrow \infty} 1 - \frac{E[\text{rank}(L)]}{QMC(G, N)}$ for all $\epsilon > 0$. So, if we can show that $Pr(\epsilon, G) \rightarrow 0$ as $\epsilon \rightarrow 0^+$, then this establishes theorem 1. For a given graph G , let $Pr(G)$ denote the limit of $Pr(\epsilon, G)$ as $\epsilon \rightarrow 0^+$.

We now prove that $Pr(G) = 0$ using induction on the number of vertices in the graph. The base case, a graph with no vertices, obviously has $Pr(G) = 0$ since all edges must be identity edges. Otherwise, given a general graph, we will either find a min cut which cuts it into two graphs with fewer vertices and apply the inductive assumption, or, if no such cut exists, we will prove that $Pr(G) = 0$ by estimating moments.

We need the following:

Lemma 9. *Let A be an N_1 -by- N_2 matrix and B be an N_2 -by- N_3 matrix. Assume that A has at least r_A singular values which are greater than or equal to ϵ_A for some ϵ_A . Assume that B has at least r_B singular values which are greater than or equal to ϵ_B for some ϵ_B . Then, AB has at least $r_A + r_B - N_2$ singular values which are greater than or equal to $\epsilon_A \epsilon_B$.*

Proof. Let P_B project onto the eigenspace of $B^\dagger B$ with eigenvalue greater than or equal to ϵ_B^2 . By assumption, P_B has rank at least r_B . We have

$$A^\dagger B^\dagger B A \geq \epsilon_B^2 A^\dagger P_B A.$$

The non-zero eigenvalues of $A^\dagger P_B A$ are the same as the nonzero eigenvalues of $P_B A A^\dagger P_B$. Let P_A project onto the eigenspace of $A A^\dagger$ with eigenvalue greater than or equal to ϵ_A^2 so that

$$P_B A A^\dagger P_B \geq \epsilon_A^2 P_B P_A P_B.$$

By assumption, P_B has rank at least r_B . The operator $P_B P_A P_B$ must have at least $r_A + r_B - N_2$ eigenvalues equal to 1 (this can be shown by using Jordan's lemma to bring both projectors P_A, P_B to a block 2-by-2 form). So, $P_B A A^\dagger P_B$ has at least $r_A + r_B - N_2$ eigenvalues greater than or equal to ϵ_A^2 and so $A^\dagger B^\dagger B A$ has at least $r_A + r_B - N_2$ eigenvalues greater than or equal to $\epsilon_A^2 \epsilon_B^2$. \square

Consider a tensor network \mathcal{N} and corresponding graph G and linear operator L . Let C be a min cut. This cut splits the tensor network into two tensor networks $\mathcal{N}_1, \mathcal{N}_2$ with corresponding linear operators L_1, L_2 so that $L = L_2 L_1$ and splits G into two graphs, G_1, G_2 . Recall that

$$K = \left(N^{|E| - MC(G)} \right)^{-1/2} L$$




FIG. VI.1: A min cut of the network shown in Fig. I.1. Thin curving vertical line shows the cut.

and similarly

$$K_1 = \left(N^{|E_1| - MC(G_1)} \right)^{-1/2} L_1, \quad K_2 = \left(N^{|E_2| - MC(G_2)} \right)^{-1/2} L_2.$$

Since $MC(G) = MC(G_1) = MC(G_2)$ and $|E_1| + |E_2| = |E| + MC(G)$ (this holds because the $MC(G)$ edges in the cut set become both input edges of G_2 and output edges of G_1), $K = K_1 K_2$. If both G_1, G_2 have at least 1 vertex, then G_1, G_2 both have fewer vertices than G and so by the inductive assumption, $Pr(G_1) = Pr(G_2) = 0$.

Lemma 10. *In this case (i.e., where $Pr(G_1) = Pr(G_2) = 0$ and where the cut splitting G into G_1, G_2 is a min cut), $Pr(G) = 0$.*

Proof. Since $Pr(G_1) = 0$, for any $\delta > 0$, there is an $\epsilon_1 > 0$ such that $Pr(\epsilon_1, G_1) \leq \delta/2$. Since $\int f(\epsilon_1) d\mu_N$ converges to $Pr(\epsilon_1, G)$, there is an N_0 such that for all $N \geq N_0$, the difference $|\int f(\epsilon_1) d\mu_N - Pr(\epsilon_1, G)|$ is bounded by $\delta/2$. So, there is an N_0 such that for all $N \geq N_0$, the probability that a randomly chosen singular value of L_1 for a random choice of tensors will be smaller than ϵ_1 is bounded by δ . So, the probability that, for a random choice of tensors in \mathcal{N}_1 , the operator L_1 will have more than $\sqrt{\delta} QMC(G_1, N)$ singular values smaller than ϵ_1 is bounded by $\sqrt{\delta}$ for all sufficiently large N .

Since $Pr(G_2) = 0$ also, there is an $\epsilon_2 > 0$ such that the probability that, for a random choice of tensors in \mathcal{N}_2 , the operator L_2 will have more than $\sqrt{\delta} QMC(G, N)$ singular values smaller than ϵ_2 is bounded by $\sqrt{\delta}$ for all sufficiently large N .

So, with probability at least $1 - 2\sqrt{\delta}$, L_1 has at least $(1 - \sqrt{\delta}) QMC(G, N)$ singular values larger than ϵ_1 and L_2 has at least $(1 - \sqrt{\delta}) QMC(G, N)$ singular values larger than ϵ_2 . So by lemma 9, with probability at least $1 - 2\sqrt{\delta}$, L has at least $(1 - 2\sqrt{\delta}) QMC(G, N)$ singular values larger than $\epsilon_1 \epsilon_2$. Since for any $\delta > 0$ such $\epsilon_1, \epsilon_2 > 0$ exist, the result follows. \square

Note that when we cut a graph, the new graph might have some open edges which connect to *no* vertices inside the graph. Such edges are open edges with one open end in S and one in T . For example, the graph in Fig. I.1 can be cut into two graphs, as shown in Fig. VI.1. We call these edges “identity edge”, because the linear operator for such a graph is the identity operator (on the degree of freedom on that edge) tensored with the linear operator for the rest of the graph. Such identity edges count as a single edge when determining $|E|$, and Eq. (I.4) still holds for a network with identity edges; one can verify this by noting that adding an identity edge multiplies the trace of any moment of $L^\dagger L$ by N and it increases both $MC(G)$ and $|E|$ by 1, leaving $|E| - MC(G)$ unchanged.

Suppose instead that G has no min cuts other than possibly the cuts $\bar{S} = S, \bar{T} = \bar{V} \setminus S$ or $\bar{S} = \bar{V} \setminus T, \bar{T} = T$; that is, the only min cuts will cut G into two graphs, one of which has no vertices. We consider three cases: (i) $|S| = MC(G) < |T|$; (ii) $|T| = MC(G) < |S|$; (iii) $|S| = |T| = MC(G)$. Case (i) and (ii) can be related by interchanging L and L^\dagger so we only do cases (i,iii) and prove in both cases that $Pr(G) = 0$.

For both cases, we will use a method of defining a new network by removing vertices from a network.

Definition 9. *Consider a tensor network \mathcal{N} . We define a new network by “removing a vertex v as input from the network” as follows. The vertex v is removed. The input edges previously attached to v are also removed. The edges attached to v which were not input edges are not removed; if they were not open edges, then they become input edges and the end attached to v becomes the open end of the edge. If they were output edges, then they become identity edges.*

Similarly, we define “removing a vertex v as output from the network”; this definition is the same as above except “input” and “output” are interchanged everywhere.

In case (i), we now show that the constant $c(G, k)$ in Eq. (I.4) is equal to 1. Thus, in this case, the limiting distribution μ is a δ -function at 1 and so $Pr(G) = 0$.

We first need the following lemma which we will also use in case (iii). Remark: even though we are proving a combinatoric result (the value of a certain constant counting pairings), we do this in a linear algebraic fashion, by estimating the trace of a certain product of operators.




FIG. VI.2: An example network for $\text{tr}((L^\dagger L)^2)$, for a graph with $|S| = MC(G) = 1$ and $|T| = 4$. For a direct pairing for which one does not pair $(v; \sigma)$ with $(v; \sigma - 1)$ for any odd σ , necessarily the pairing is the same as that constructed in lemma 4.

Lemma 11. Let $\mathcal{N}(1), \mathcal{N}(2), \dots, \mathcal{N}(l)$ be tensor networks with open edges, with $\mathcal{N}(i)$ having the same number of output edges as $\mathcal{N}(i+1)$ has input edges (identifying $\mathcal{N}_c(l+1)$ with $\mathcal{N}_c(1)$). Let L_1, L_2, \dots, L_l be the corresponding linear operators and let G_1, \dots be the corresponding graphs, with edge sets $E(1), \dots$. Consider the tensor network \mathcal{N}_c which computes the trace $\text{tr}(L_l L_{l-1} \dots L_1)$. This tensor network has

$$C_{\max} \leq \sum_{i=1}^l \left(|E(i)| - MC(G(i)) \right) + \min_i MC(G(i)). \quad (\text{VI.2})$$

Proof. By corollary 2, for all $n, \epsilon_i > 0$, for all sufficiently large N , there is a c_i such that the probability that the largest singular value of $L(i)$ is $\geq x N^{|E(i)| - MC(G(i)) + \epsilon_i}$ is bounded by $c_i x^{-n}$.

The largest singular value of $L_l L_{l-1} \dots L_1$ is bounded by the product over i of the largest singular values of L_i , so for any m there is a constant c such that the probability that the largest singular value of $L_l L_{l-1} \dots L_1$ is $\geq z N^{\sum_i (|E(i)| - MC(G(i)) + \epsilon_i)}$ is bounded by $c z^{-m}$. (To show this bound, one must pick each n sufficiently large above) Choosing $\epsilon_i = \epsilon/l$, it follows that for all $\epsilon > 0$, for all m , there is a constant c such that the probability the largest singular value of $L_l L_{l-1} \dots L_1$ is $\geq z N^{\sum_i (|E(i)| - MC(G(i)) + \epsilon)}$ is bounded by $c z^{-m}$.

The operator $L_l L_{l-1} \dots L_1$ has rank at most $\min_i MC(G(i))$ so the probability that the trace $\text{tr}(L_l L_{l-1} \dots L_1)$ is $\geq z N^{\sum_i (|E(i)| - MC(G(i)) + \min_i MC(G(i)) + \epsilon)}$ is bounded by $c z^{-m}$. Choosing m sufficiently large, this implies that $E[\text{tr}(L_l L_{l-1} \dots L_1)]$ is $\mathcal{O}(N^{\sum_i (|E(i)| - MC(G(i)) + \min_i MC(G(i)) + \epsilon)})$ for all $\epsilon > 0$. Choosing $\epsilon < 1$, this implies the result. \square

Lemma 12. For a graph G in case (i), $c(G, k) = 1$.

Proof. Let $\mathcal{N}_c = \mathcal{N}_c(\text{tr}((L^\dagger L)^k))$. We show that $n_{\max} = 1$ for all k . We only need to consider the case $k > 1$ (for $k = 1$, lemma 3 shows $n_{\max} = 1$). Lemma 4 constructs a direct pairing π with $C(\pi) = C_{\max} = k|E| - (k-1)MC(G)$. There is only one min cut of the graph, so this lemma constructs only one such direct pairing. Let π' be a direct pairing different from that π . Since $\pi' \neq \pi$, there must be some vertex $(v; \sigma)$ for odd σ which is attached to an input edge such that we pair $(v; \sigma)$ with $(v; \sigma - 1)$. For example, see Fig. VI.2.

Consider the sum over all pairings π'' which pair $(v; \sigma)$ with $(v; \sigma - 1)$, weighted by $N^{C(\pi'')}$. This is the expectation value of the one-step direct subnetwork made by pairing $(v; \sigma)$ with $(v; \sigma - 1)$. Let L be the linear operator defined above, and let M be the linear operator corresponding to the network with vertex v removed as input.

Then, the contraction of this one-step direct subnetwork is equal to $\text{tr}((L^\dagger L)^{k-1} M^\dagger M)$. The graph G_M corresponding to M has $MC(G_M) \geq MC(G) + 1$ (if not, one has a min cut of G with partition $\bar{V} = \bar{S} \cup \bar{T}$ with $\bar{S} = S \cup \{v\}$). Let there be N_E edges connecting $(v; \sigma)$ with $(v; \sigma - 1)$ so that G_M has $|E| - N_E$ edges. Hence by lemma 11, the network \mathcal{N}'_c has $C_{\max} \leq k|E| - (k-1)MC(G) - N_E - 1$ and so, all pairings π' of \mathcal{N}_c which pair $(v; \sigma)$ with $(v; \sigma - 1)$ have $C(\pi') \leq k|E| - (k-1)MC(G) - 1$. \square

In case (iii), we show that the constant $c(G, k)$ in Eq. (I.4) is independent of the particular graph chosen so long as the graph has at least 1 vertex. In particular, the constant is the same as that if we consider a graph with $|S| = |T| = MC(G) = 1$, with one vertex of degree 2 and $|E| = 2$, with both edges open. This case is the well-known case of random matrix theory studying the eigenvalues of $L^\dagger L$ where L is a random N -by- N matrix with independent Gaussian entries; the matrix L is drawn from the chiral Gaussian unitary ensemble. In this case, the limiting distribution is known⁵ and is smooth near 0 so again $Pr(G) = 0$.

Lemma 13. For a graph G in case (iii), $c(G, k)$ is independent of G so long as G has at least 1 vertex.

Proof. Let $\mathcal{N}_c = \mathcal{N}_c(\text{tr}((L^\dagger L)^k))$. Consider the sum of all pairings π in which for some odd σ , for some vertex v attached to an input edge, $(v; \sigma)$ is paired with $(v; \sigma - 1)$ and for some vertex w attached to an output edge $(w; \sigma)$ is paired with $(w; \sigma + 1)$. This is the expectation value of the direct subnetwork \mathcal{N}'_c made by pairing $(v; \sigma)$ with




FIG. VI.3: A network with $|S| = |T| = MC(G) = 2$ and with no min cuts other than those which divide G into two graphs, one of which has no vertices. Input edges are on left and output edges are on right.

$(v; \sigma - 1)$ and $(w; \sigma)$ with $(w; \sigma + 1)$. Let L be the linear operator defined above, and let M_1 be the linear operator corresponding to the network with vertex v removed as output and let M_2 be the linear operator corresponding to the network with vertices v removed as input and w removed as output and let M_3 be the linear operator corresponding to the network with vertex w removed as input. Then, the contraction of \mathcal{N}'_c equals $\text{tr}((L^\dagger L)^{k-2} M_1^\dagger M_2 M_3^\dagger L)$. Let $G_{M_1}, G_{M_2}, G_{M_3}$ be the graphs corresponding to M_1, M_2, M_3 , respectively. We have $MC(G_{M_2}) \geq MC(G) + 1$ (if not, then there is a set of $MC(G)$ edges that one can remove from G_{M_2} to disconnect the graph; removing this same set of edges from G will disconnect G giving a min cut which cuts G into two graphs, each of which has at least one vertex).

Let there be $N_E(v)$ edges connecting $(v; \sigma)$ with $(v; \sigma - 1)$ and $N_E(w)$ edges connecting $(v; \sigma)$ with $(v; \sigma + 1)$, so G_{M_1} has $|E| - N_E(v)$ edges, G_{M_2} has $|E| - N_E(v) - N_E(w)$ edges, and G_{M_3} has $|E| - N_E(w)$ edges. Hence, by lemma 11, \mathcal{N}'_c has $C_{\max} \leq k|E| - (k-1)MC(G) - N_E(v) - N_E(w) - 1$, and so all such pairings π of \mathcal{N}_c have $C(\pi) \leq k|E| - (k-1)MC(G) - 1$.

The above proof was for σ odd, but it works with σ even also, if one takes the adjoint of all linear operators L, M_1, M_2, M_3 .

So, for a maximal pairing, there is no σ such that for some vertex v attached to an input edge, $(v; \sigma)$ is paired with $(v; \sigma - 1)$ and for some vertex w attached to an output edge $(w; \sigma)$ is paired with $(w; \sigma + 1)$.

For a direct pairing, there must be some σ such that for some vertex v attached to an input edge, $(v; \sigma)$ is paired with $(v; \sigma - 1)$ or for some vertex w attached to an output edge $(w; \sigma)$ is paired with $(w; \sigma + 1)$. So, for that σ , either for *all* v , $(v; \sigma)$ is paired with $(v; \sigma - 1)$ or for *all* v , $(v; \sigma)$ is paired with $(v; \sigma + 1)$. Suppose, without loss of generality, the first case holds, so that for all v , $(v; \sigma)$ is paired with $(v; \sigma - 1)$ for that σ . Define a direct subnetwork of \mathcal{N}_c by pairing such vertices v in that way; this direct subnetwork is precisely the network $\mathcal{N}_c(\text{tr}((L^\dagger L)^{k-1}))$, so one can apply the result above and for some σ , either for *all* v , $(v; \sigma)$ is paired with $(v; \sigma - 1)$ or for *all* v , $(v; \sigma)$ is paired with $(v; \sigma + 1)$. Iterating this k times, one is left with a network with no vertices. The possible choices (which σ and whether one pairs with $\sigma - 1$ or $\sigma + 1$) are in one-to-one correspondence with rainbow diagrams and do not depend upon G . \square

To better understand this case $|S| = |T| = MC(G)$, consider the network shown in Fig. VI.3 as an example of such a network. Imagine generalizing the networks that we consider, so that the 4 open edges have capacity N but the vertical edge in the figure has some capacity $N' \leq N$. In the case $N' = 1$, the tensor network factors into a product of two tensor networks; the linear operator L factors as $L = L_1 \otimes L_2$, where L_1 maps the upper input edge to the upper output edge and L_2 maps the lower input edge to the lower output edge (upper and lower refer to the position of the edge in the figure). Then, the singular value spectrum of L is the product of the singular value spectrums of L_1, L_2 and the limiting distribution of the singular value spectrum of L is *not* the same as in the chiral Gaussian unitary ensemble. (The case with $N' = 1$ is equivalent to a case with degree $d = 2$ but where there is a min cut which splits the graph into two networks each with only 1 vertex; i.e., simply redraw the graph by removing the vertical edge with $N = 1$.) For general N' , the entries of L can be obtained by summing N' independent matrices of the form $L_1 \otimes L_2$. As N' increases, the correlations between entries of L are reduced until eventually when N' is large enough, the singular value spectrum of L is the same as in the chiral Gaussian unitary ensemble.

VII. VARIANCE

Here we collect some results on variance of moments.

The first result bounds the variance

Lemma 14. *For any linear operator L obtained from a tensor network \mathcal{N}_c with graph G with $|E|$ edges, for any*

positive integer c

$$E[\text{tr}((L^\dagger L)^k)^c] - E[\text{tr}((L^\dagger L)^k)]^c = \mathcal{O}(N^{c(k|E|-(k_i-1)MC(G))-1}), \quad (\text{VII.1})$$

and

$$E_{\text{ind}}[\text{tr}((L^\dagger L)^k)^c] - E_{\text{ind}}[\text{tr}((L^\dagger L)^k)]^c = \mathcal{O}(N^{2(k|E|-(k_i-1)MC(G))-1}), \quad (\text{VII.2})$$

Proof. Let $M = L \otimes L \dots$, with a total of c copies of L in the tensor product so that $\text{tr}((L^\dagger L)^k)^c = \text{tr}((M^\dagger M)^k)$. Then, M is the linear operator associated to the tensor network \mathcal{N}'_c which is the product $\mathcal{N}_c \cdot \mathcal{N}_c \dots$ and has graph G' with min cut $MC(G') = cMC(G)$. We can label vertices in G' by a pair $\{v; d\}$ for $d = 1, 2, \dots, c$ where d labels the two different tensor factors in M (we use different notation $\{v; f\}$ to distinguish this from the notation $(v; \sigma)$ used above). Then, vertices in \mathcal{N}'_c are labelled by triples $(\{v; d\}, \sigma)$. We can use lemma 8 to show that the pairings of \mathcal{N}'_c with maximal number of closed loops are direct pairings. Every direct pairing of \mathcal{N}'_c only pairs $(\{v; d\}, \sigma)$ with $(\{v; d\}, \tau)$, so that the direct pairings of \mathcal{N}'_c are in one-to-one correspondence to the product of the direct pairings of \mathcal{N}_c with itself. \square

This has the corollary using the case $c = 2$:

Corollary 3. *With high probability, for either choice of ensemble, $\text{av}((K^\dagger K)^k)$ is within $o(1)$ of its expectation value.*

To study expectation values of other products, we define a “product of pairings”.

Definition 10. *Let $\mathcal{N}_c(1), \mathcal{N}_c(2), \dots, \mathcal{N}_c(a)$ be a sequence of closed tensor networks, $a \geq 1$. Let $\mathcal{N}_c = \mathcal{N}_c(1) \dots \mathcal{N}_c(a)$ be the product of these networks and label each vertex in the product network by a pair $[v; i]$, for $1 \leq i \leq a$, where v labels a vertex in tensor network $\mathcal{N}_c(i)$. Let $\pi(1), \pi(2), \dots, \pi(n)$ be pairings of these networks. Then, define the pairing π which is a product of these pairings to be the pairing which pairs $[v; i]$ with $[w; i]$ if $\pi(i)$ pairs v with w .*

Note that $C(\pi) = \sum_i C(\pi(i))$.

Then, once we lower bound a trace, we have an immediate lower bound for products of traces:

Lemma 15. *Let $\mathcal{N}_c(1), \mathcal{N}_c(2), \dots, \mathcal{N}_c(a)$ be a sequence of closed tensor networks, $a \geq 1$. Let $C_{\max}(i)$ be the maximum number of loops in a pairing of $\mathcal{N}_c(i)$ and let $n_{\max}(i)$ be the number of distinct pairings π of $\mathcal{N}_c(i)$ with $C(\pi) = C_{\max}(i)$. Then, for the network \mathcal{N}_c defined to be the product network $\mathcal{N}_c(1) \cdot \mathcal{N}_c(2) \cdot \dots \cdot \mathcal{N}_c(a)$ we have*

$$C_{\max} \geq \sum_i C_{\max}(i), \quad (\text{VII.3})$$

and if $C_{\max} = \sum_i C_{\max}(i)$ then

$$n_{\max} \geq \prod_i n_{\max}(i). \quad (\text{VII.4})$$

Further,

$$E[\mathcal{N}_c] \geq E[\mathcal{N}_c(1)] \cdot E[\mathcal{N}_c(2)] \cdot \dots \cdot E[\mathcal{N}_c(a)], \quad (\text{VII.5})$$

where $E[\mathcal{N}_c]$ denotes the expectation value of the contraction of the network.

Proof. Let $\pi_i(j)$, for $1 \leq j \leq n_{\max}(i)$ label the pairings of $\mathcal{N}_c(i)$ with $C(\pi) = C_{\max}(i)$. Then, for each sequence j_1, \dots, j_a , consider the product pairing of \mathcal{N}_c . This shows Eqs. (VII.3, VII.4). To show Eq. (VII.5), note that each expectation value $E[\mathcal{N}_c(i)]$ is a sum over pairings of that network, weighted by the number of closed loops; the sum over product pairings of \mathcal{N}_c , weighted by the number of closed loops, is equal to $E[\mathcal{N}_c(1)] \cdot \dots \cdot E[\mathcal{N}_c(a)]$. \square

We now show:

Lemma 16. *The expectation value of any product of traces is close to the product of the averages:*

$$E[\text{tr}((L^\dagger L)^{k_1}) \text{tr}((L^\dagger L)^{k_2}) \dots] = (1 + \mathcal{O}(1/N)) \cdot E[\text{tr}((L^\dagger L)^{k_1})] E[\text{tr}((L^\dagger L)^{k_2})] \dots \quad (\text{VII.6})$$

Proof. By lemma 15, $E[\text{tr}((L^\dagger L)^{k_1})\text{tr}((L^\dagger L)^{k_2})] \dots \geq (1 - \mathcal{O}(1/N)) \cdot E[\text{tr}((L^\dagger L)^{k_1})]E[\text{tr}((L^\dagger L)^{k_2})] \dots$. So, we just need to upper bound the left-hand side of Eq. (VII.6).

Consider the case of a product of only two traces. Then, by Cauchy-Schwarz, $E[\text{tr}((L^\dagger L)^{k_1})\text{tr}((L^\dagger L)^{k_2})] \leq \sqrt{E[\text{tr}((L^\dagger L)^{k_1})^2]E[\text{tr}((L^\dagger L)^{k_2})^2]}$. By lemma 14, and Eq. (I.4),

$$\sqrt{E[\text{tr}((L^\dagger L)^k)^2]} \leq (1 + \mathcal{O}(1/N)) \cdot E[\text{tr}((L^\dagger L)^k)]. \quad (\text{VII.7})$$

So,

$$E[\text{tr}((L^\dagger L)^{k_1})\text{tr}((L^\dagger L)^{k_2})] \leq (1 + \mathcal{O}(1/N)) \cdot E[\text{tr}((L^\dagger L)^{k_1})]E[\text{tr}((L^\dagger L)^{k_2})]. \quad (\text{VII.8})$$

So, the claim follows for the case of two traces.

Now, consider a case with a traces for some $a > 2$. By Cauchy-Schwarz,

$$E[\text{tr}((L^\dagger L)^{k_1})\text{tr}((L^\dagger L)^{k_2}) \dots \text{tr}((L^\dagger L)^{k_a})] \leq \sqrt{E[\text{tr}((L^\dagger L)^{k_1})^2]} \sqrt{E[\text{tr}((L^\dagger L)^{k_2})^2] \dots \text{tr}((L^\dagger L)^{k_a})^2]}. \quad (\text{VII.9})$$

Applying Cauchy-Schwarz again to the second term in this product on the right-hand side, and continuing in this fashion, we find that

$$E[\text{tr}((L^\dagger L)^{k_1})\text{tr}((L^\dagger L)^{k_2}) \dots \text{tr}((L^\dagger L)^{k_a})] \leq \dots E[\text{tr}((L^\dagger L)^{k_1})^2]^{1/2} E[\text{tr}((L^\dagger L)^{k_2})^4]^{1/4} E[\text{tr}((L^\dagger L)^{k_3})^8]^{1/8} \dots \quad (\text{VII.10})$$

By lemma 14, and Eq. (I.4),

$$E[\text{tr}((L^\dagger L)^{k_1})\text{tr}((L^\dagger L)^{k_2}) \dots] \leq (1 + \mathcal{O}(1/N)) \cdot E[\text{tr}((L^\dagger L)^{k_1})]E[\text{tr}((L^\dagger L)^{k_2})] \dots \quad (\text{VII.11})$$

□

VIII. NUMERICS

We performed a numerical study of the network shown in Fig. VIII.1. This network was used in Ref. 2 as an example of a network for which $QMF < QMC$ for $N = 2$. Since this network has $|S| = |T| = MC(G) = 2$ and there is no min cut that cuts the graph into two graphs which each have at least one vertex, the results above predict that for large N , the limiting distribution of singular values is, up to dividing by $\sqrt{N^{|E|-MC}} = N^3$, the same as that for random matrix theory in the chiral Gaussian unitary ensemble. In Fig. VIII.2 we show the singular values both for a randomly chosen example of this ensemble for $N = 20$ and for a randomly chosen chiral Gaussian unitary matrix of size N^2 -by- N^2 . In Fig. VIII.3, we show the same for $N = 40$. One can see the convergence.

We have also considered the smallest singular values to obtain numerical evidence for the rank of L . Studying a few random samples for each N with $2 \leq N \leq 40$, we found that there was exactly one very small singular value (smaller than 10^{-14}) for $N = 2, 3 \pmod{4}$ and no very small singular values otherwise. In both cases, all the other singular values were larger than $6 * 10^{-4}$. While this is only numerical evidence, it supports a conjecture that for this network, $QMF(G, N, O) = QMC(G, N) - 1$ if $N = 2, 3 \pmod{4}$ and $QMF(G, N, O) = QMC(G, N)$ otherwise.

The quantum max-flow/min-cut conjecture of Ref. 1 was shown to be false in Ref. 2. If the conjecture above for the network shown in Fig. VIII.1 is true, then it is not even the case that “for all G, O , for all sufficiently large N , $QMF(G, N, O) = QMC(G, N)$ ”. We may still hope for the weaker conjecture that “for all G, O , for infinitely many N , $QMF(G, N, O) = QMC(G, N)$ ”. Indeed, to show this weaker conjecture, it suffices to show that “for all G, O , for some $N_0 > 1$, $QMF(G, N_0, O) = QMC(G, N_0)$ ” as then $QMF(G, N_0^k, O) = QMC(G, N_0^k)$ for all integer $k \geq 0$ due to the following:

Lemma 17. *For all G , $QMF(G, N_1 N_2, O) \geq QMF(G, N_1, O)QMF(G, N_2, O)$.*

Proof. Let $\mathcal{T}_1, \mathcal{T}_2$ be tensors whose indices range from 1 to N_1 or 1 to N_2 respectively, such that the network with tensor \mathcal{T}_1 gives a linear operator L_1 with rank $QMF(G, N_1, O)$ and with tensor \mathcal{T}_2 gives a linear operator L_2 with rank $QMF(G, N_2, O)$. Then, $\mathcal{T} = \mathcal{T}_1 \otimes \mathcal{T}_2$ is a tensor whose indices range from 1 to $N_1 N_2$ and gives a linear operator $L = L_1 \otimes L_2$ with rank $QMF(G, N_1, O)QMF(G, N_2, O)$. Here, the product $\mathcal{T}_1 \otimes \mathcal{T}_2$ is defined as follows: label the each index of \mathcal{T} by a pair (i, j) . Write the entries of \mathcal{T} as $\mathcal{T}_{(i_1, j_1), (i_2, j_2), \dots, (i_d, j_d)}$. Set $\mathcal{T}_{(i_1, j_1), (i_2, j_2), \dots, (i_d, j_d)} = (\mathcal{T}_1)_{i_1, i_2, \dots, i_d} (\mathcal{T}_2)_{j_1, j_2, \dots, j_d}$. □




FIG. VIII.1: A tensor network from Ref. 2 that we studied numerically. Numbers indicate local ordering.




FIG. VIII.2: Plot of the value of the singular values (divided by N^3) for one sample of the tensor network shown in Fig. VIII.1 for $N = 20$ shown in blue and for a chiral Gaussian unitary matrix of size N^2 -by- N^2 shown in green. The largest singular value shown on the blue curve is larger than the largest on the green curve.

¹ D. Calegari, M. Freedman, and K. Walker, “Positivity of the universal pairing in 3 dimensions”, Jour. Amer. Math. Soc. **23**, no. 1, 107-188 (2010).

² S. X. Cui, M. H. Freedman, O. Sattath, R. Stong, and G. Minton, “Quantum Max-flow/Min-cut”, arXiv:1508.04644.

³ For a sequence of distributions where the moments converge to a limit $c(G, k)$ obeying Carleman’s condition $\sum_{k=1}^{\infty} c(G, 2k)^{\frac{1}{2k}} = \infty$, the distributions converge weakly to a limiting distribution. See <http://mathoverflow.net/questions/230794/converging-to-moments-obeying-carlemans-condition>. Sketch: since the second moment is bounded, the sequence is tight; by Prokhorov’s theorem, there is a subsequence which converges to a limit μ ; the moments of the limiting distribution μ are the limits of the moments (to show that μ has the correct k -th moment, use a bound on a higher moment to establish uniform integrability) so by Carleman’s condition, all subsequences which converge must converge to the same limit; so by Prokhorov’s theorem, the sequence converges without needing to pass to a subsequence. In this particular case, since the $2k$ -th moments of μ_N^{ind} (or μ_N) are bounded by $c_1 \cdot c_2^k$ for some constants c_1, c_2 , it is simpler. To show that $\int f(x) d\mu_N^{\text{ind}}(x)$ has a limit as $N \rightarrow \infty$ for bounded Lipschitz functions $f(x)$, approximate $f(x)$ on an interval $[-c, +c]$ by a polynomial $p(x)$ for any $c > c_2$ and use a bound on higher moments to bound the integral $\int_{|x|>c} p(x) d\mu_N^{\text{ind}}(x)$.

⁴ “Rainbow diagrams” for random matrix theory are another term for “planar diagrams”. This is a class of diagrams that appear in computing expectation values of traces of powers of a random matrix. See, for example, A. Zee, *Quantum Field*




FIG. VIII.3: Plot of the value of the singular values (divided by N^3) for one sample of the tensor network shown in Fig. VIII.1 for $N = 40$ shown in blue and for a chiral Gaussian unitary matrix of size N^2 -by- N^2 shown in green. The largest singular value shown on the blue curve is slightly larger than the largest on the green curve but the discrepancy is smaller than in the case $N = 20$.

Theory in a Nutshell, Second Edition, pp 396-400, Princeton University Press (Princeton, NJ 2010).

⁵ E.V. Shuryak and J.J.M. Verbaarschot, Nucl. Phys. A **560**, 306 (1993).

⁶ J.J.M. Verbaarschot, Phys. Rev. Lett. **72**, 2531 (1994); Phys. Lett. B **329**, 351 (1994).

⁷ Since the limiting distribution in the Gaussian orthogonal ensemble is the Wigner semi-circle, which is bounded, one can deduce that the expected trace of the k -th moment of a matrix drawn from this ensemble is at most exponential in k . See also E. Brezin, C. Itzykson, G. Parisi, and J. B. Zuber, "Planar Diagrams", Commun. Math. Phys. **59**, 35 (1978). Alternatively, one can reduce the problem of counting rainbow diagrams in this case to the problem of counting parenthesizations, which are bounded by an exponential.

⁸ P. Elias, A. Feinstein, and C. E. Shannon, "A note on the maximum flow through a network", Information Theory, IRE Transactions on, **2**(4), 117-119 (1956).

⁹ L. R. Ford and D. R. Fulkerson, "Maximal flow through a network", Canadian journal of Mathematics, **8**(3), 399-404 (1956).